

# **SAN Host Utilities 3.3.1 for iSCSI on Linux®**

## **Release Notes**

NetApp, Inc.  
495 East Java Drive  
Sunnyvale, CA 94089 USA  
Telephone: +1 (408) 822-6000  
Fax: +1 (408) 822-4501  
Support telephone: +1 (888) 4-NETAPP  
Documentation comments: [doccomments@netapp.com](mailto:doccomments@netapp.com)  
Information Web: <http://www.netapp.com>

Part number: 215-04587\_A0  
February 2009

# Copyright and trademark information

---

## Copyright information

Copyright © 1994-2009 NetApp, Inc. All rights reserved. Printed in the U.S.A.

The third-party copyright statements and licenses that apply to open source software components distributed with the product are contained in the section “Third-Party Copyrights and Licenses.”

No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S.A. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

## Trademark information

NetApp, the Network Appliance logo, the bolt design, NetApp—the Network Appliance Company, Cryptainer, Cryptoshred, DataFabric, DataFort, Data ONTAP, Decru, FAServer, FilerView, FlexClone, FlexVol, Manage ONTAP, MultiStore, NearStore, NetCache, NOW NetApp on the Web, SANscreen, SecureShare, SnapDrive, SnapLock, SnapManager, SnapMirror, SnapMover, SnapRestore, SnapValidator, SnapVault, Spinnaker Networks, SpinCluster, SpinFS, SpinHA, SpinMove, SpinServer, StoreVault, SyncMirror, Topio, VFM, VFM (Virtual File Manager), and WAFL are registered trademarks of NetApp, Inc. in the U.S.A. and/or other countries. gFiler, Network Appliance, SnapCopy, Snapshot, and The evolution of storage are trademarks of NetApp, Inc. in the U.S.A. and/or other countries and registered trademarks in some other countries. The NetApp arch logo; the StoreVault logo; ApplianceWatch; BareMetal; Camera-to-Viewer; ComplianceClock; ComplianceJournal; ContentDirector; ContentFabric; EdgeFiler; FlexShare; FPolicy; Go Further, Faster; HyperSAN; InfoFabric; Lifetime Key Management, LockVault; NOW; ONTAPI; OpenKey, RAID-DP; ReplicatorX; RoboCache; RoboFiler; SecureAdmin; Serving Data by

Design; SharedStorage; Simplicore; Simulate ONTAP; Smart SAN; SnapCache; SnapDirector; SnapFilter; SnapMigrator; SnapSuite; SohoFiler; SpinMirror; SpinRestore; SpinShot; SpinStor; vFiler; Virtual File Manager; VPolicy; and Web Filer are trademarks of NetApp, Inc. in the U.S.A. and other countries. NetApp Availability Assurance and NetApp ProTech Expert are service marks of NetApp, Inc. in the U.S.A.

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. A complete and current list of other IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>.

Apple is a registered trademark and QuickTime is a trademark of Apple, Inc. in the U.S.A. and/or other countries. Microsoft is a registered trademark and Windows Media is a trademark of Microsoft Corporation in the U.S.A. and/or other countries. RealAudio, RealNetworks, RealPlayer, RealSystem, RealText, and RealVideo are registered trademarks and RealMedia, RealProxy, and SureStream are trademarks of RealNetworks, Inc. in the U.S.A. and/or other countries.

All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

NetApp, Inc. is a licensee of the CompactFlash and CF Logo trademarks. NetApp, Inc. NetCache is certified RealSystem compatible.



# Table of Contents

---

|                                                                                            |    |
|--------------------------------------------------------------------------------------------|----|
| Changes in this release . . . . .                                                          | 1  |
| Contents of Host Utilities. . . . .                                                        | 2  |
| Known problems and limitations using RHEL 5 series . . . . .                               | 4  |
| Known problems and limitations using RHEL 4 series . . . . .                               | 6  |
| Known problems and limitations using RHEL 3 series . . . . .                               | 12 |
| Upgrading the operating system from the previous version . . . . .                         | 17 |
| Known problems and limitations on SUSE Linux Enterprise<br>Server 10 SP1 and SP2 . . . . . | 18 |
| Known problems and limitations on SUSE Linux Enterprise Server 9 . . . .                   | 21 |
| Information about additional problems. . . . .                                             | 27 |
| Where to find additional information. . . . .                                              | 28 |



# Changes in this release

---

## New feature in the release

The iSCSI Linux® Host Utilities 3.3.1 Release Notes update of February 2009:

New multipath.conf parameters have been added in Red Hat Enterprise Linux 5 Update 3. Following are the parameters that are added to /etc/multipath.conf file:

- ◆ `max_fds`  
The `max_fds` parameter is now set to `max`, which sets the maximum file descriptors for the `multipathd` process. The value of this parameter must be set to a higher value than the number of paths to be configured.
- ◆ `flush_on_last_del`  
The `flush_on_last_del` parameter is set to `yes`. The `multipathd` process disables queueing when the last path to a device is deleted.
- ◆ `queue_without_daemon`  
The `queue_without_daemon` parameter is set to `no`. The `multipathd` process disables queueing when the multipath daemon is stopped.

The following example shows the required changes to the /etc/multipath.conf file on Red Hat Enterprise Linux 5 Update 3:

```
defaults
{
user_friendly_names      yes
max_fds                  max
queue_without_daemon     no
}
blacklist
{
wwid SIBM-ESXSMAT3073NC_FAAR9P5A0ADG6
wwid SIBM-ESXSMAT3073NC_FAAR9P5A0ADKP
devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st) [0-9]*"
devnode "^hd[a-z]"
devnode "^cciss!c[0-9]d[0-9]*[p[0-9]]*"
}
devices
{
device
{
vendor      "NETAPP"
product     "LUN"
path_grouping_policy multibus
}
```

```

getuid_callout      "/sbin/scsi_id -g -u -s /block/%n"
prio_callout        "/sbin/mpath_prio_ontap /dev/%n"
features            "1 queue_if_no_path"
path_checker        directio
failback            immediate
flush_on_last_del   yes
}
}

```

## Documentation changes

Updated previous documentation-related issues.

## Supported packages

Use the version of the packages that are shipped along with the specific operating system release, unless explicitly mentioned.

# Contents of Host Utilities

---

## Software

Host Utilities includes a set of software programs to help you use your Linux host with storage systems. You can download this software from the NOW (Netpp on the Web)<sup>®</sup> site after accepting the NetApp Customer Software License.

The software is distributed as a single file (**netapp\_linux\_host\_utilities\_3\_3.tar.gz**), which you can install on your Linux host. The installation instructions are provided in the *Setup Guide* and on the NOW download page.

The same software distribution is shared by a number of Host Utilities. Therefore, it contains some software that is not used for Linux iSCSI. The following table lists each file in the distribution:

| File                      | Purpose                                                                                                                                           |
|---------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------|
| install                   | Installs the Host Utilities software on a Linux host.                                                                                             |
| uninstall                 | Removes the Host Utilities software from the Linux host.                                                                                          |
| sanlun_bin_ne<br>tapp.tar | Contains the following files.                                                                                                                     |
| sanlun                    | Displays information about LUNs on connected storage systems.                                                                                     |
| filer_info                | Collects information about the storage system in a file that can be submitted to NetApp for problem analysis.                                     |
| linux_info                | Collects information about the Linux host in a file that can be submitted to NetApp for problem analysis.                                         |
| san_version               | Displays the version of the software tools installed for this Host Utilities. The version of the software tools for this Host Utilities is 3.3.1. |
| brocade_info              | Not used for iSCSI.                                                                                                                               |
| cisco_info                | Not used for iSCSI.                                                                                                                               |
| mcdata_info               | Not used for iSCSI.                                                                                                                               |
| Telnet.pm                 | Support library for *info scripts.                                                                                                                |

| File             | Purpose                                                    |
|------------------|------------------------------------------------------------|
| mpath_prio_ontap | Path priority callout program used in /etc/multipath.conf. |
| SHsupport.pm     | Support library for *info scripts.                         |
| libqlsdm.so      | Support library for sanlun.                                |

## Documentation

Host Utilities includes the following documentation.

| Document                | Purpose                                                                                                                                                                                                                                                                                                             |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Release Notes           | Lists the contents of Host Utilities, and any known problems or last minute updates.<br><br>Download from the NOW Web site.                                                                                                                                                                                         |
| Setup Guide             | Explains how to install and configure software to enable the Linux host to use LUNs on the storage system.<br><br>Download from the NOW Web site.                                                                                                                                                                   |
| Command Quick Reference | Lists the most common storage system and Linux commands used to manage iSCSI.<br><br>Download from the NOW Web site.                                                                                                                                                                                                |
| man pages               | Displays information about the Host Utilities software.<br><br>The man pages for the <b>sanlun</b> , <b>filer_info</b> , and <b>linux_info</b> commands are copied to the /opt/netapp/man/man1 directory when you install the Host Utilities software.<br><br>Included in the Host Utilities software distribution. |
| NOW pages               | Host Utilities installation instructions, information updates, and links to additional information.<br><br>Browse on the NOW Web site.                                                                                                                                                                              |

## Known problems and limitations using RHEL 5 series

---

### GFS growth is not possible on mounted file system on an iSCSI LUN

**Problem:** GFS provides the online grow feature, where the GFS file system can be extended without unmounting the file system by using the **gfs\_grow** command. This feature cannot be used on a file system created over an iSCSI and multipath device on a NetApp LUN.

---

#### Note

When GFS is used on the LVM volume, it can be extended without unmounting the file system by adding new LUNs to the volume group, and expanding the logical volume on which GFS is created.

---

Though storage systems provide dynamic resizing of the LUN, the iSCSI layer in Linux is not capable of detecting the LUN resize. The new size of the LUN gets reflected on Linux only after the iSCSI service is restarted on Linux.

**Workaround:** Complete the following steps on all the nodes.

| Step | Action                                                       |
|------|--------------------------------------------------------------|
| 1    | Unmount the file system.                                     |
| 2    | Resize the LUN on the storage system.                        |
| 3    | Run the following command:<br><b>multipath -F</b>            |
| 4    | Run the following command:<br><b>service multipath stop</b>  |
| 5    | Run the following command:<br><b>service iscsi restart</b>   |
| 6    | Run the following command:<br><b>service multipath start</b> |
| 7    | Mount the file system.                                       |

| Step | Action                                                                                                   |
|------|----------------------------------------------------------------------------------------------------------|
| 8    | Run the following command on one of the nodes:<br><b>gfs_grow</b>                                        |
| 9    | Run the following command on all nodes to verify the new size:<br><b>running gfs_grow -Tv mountpoint</b> |

For more information, see bug 215923 on Bugs Online at:

<http://now.netapp.com/NOW/cgi-bin/bol>

### Using partitions on Linux with Data ONTAP® LUNs might require alignment for best performance

**Problem:** Using default partition settings on LUNs might result in performance problems. This problem can occur if the `sectors` value for a partition is not a multiple of 8.

**Workaround:** Manually enter the disk geometry parameters when creating partitions with the **fdisk** command. For specific steps, see bug number 156121 on Bugs Online at:

<http://now.netapp.com/NOW/cgi-bin/bol>

### Host freezes intermittently if multipathd daemon is not running in SAN boot scenario

**Problem:** During controller faults, all paths go down for some time. Also, if the multipathd daemon is not running, the root volume is unavailable, which leads to a system freeze. Another scenario can be a host shutdown. If paths go down after the multipath service is stopped, the host freezes, and you should do a hard reboot.

**Workaround:** Ensure that the multipathd daemon does not stop when the root volume is on the iSCSI multipathed device. If the multipathd daemon is up and running, the `queue_if_no_path` feature holds I/O on the root volume till some paths get reinstated. After one path gets reinstated, the root volume is accessible.

To ensure that the daemon does not stop, you have to modify the multipathd init script. You can modify the script by applying the patch, which is available at the following Web site:

<https://bugzilla.redhat.com/attachment.cgi?id=314924>

For more information, see NetApp bug 295559 on Bugs Online at:

<http://now.netapp.com/NOW/cgi-bin/bol>

## RHEL 5 Update 2 host intermittently stops while booting in SAN boot scenario

**Problem:** There is a problem with the iscsid and multipathd sequence during the boot process. Because the iscsid service is started before the multipathd service, there is a small window when the host sometimes loses connection with the iSCSI root volume, which leads to a system freeze.

**Workaround:** Configure the system after postinstallation by completing the following steps.

| Step | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | During the first boot, switch to the single user mode by modifying the grub parameter.                                                                                                                                                                                                                                                                                                                                                                                            |
| 2    | Add <b>single</b> to the kernel parameters.<br><br>Example: <b>kernel /vmlinuz-2.6.18-92.el5 ro root=/dev/VolGroup00/LogVol100 rhgb quiet single</b>                                                                                                                                                                                                                                                                                                                              |
| 3    | Create the /etc/multipath.conf file as described in “Configuring the /etc/multipath.conf file” in the <i>Setup Guide</i> .                                                                                                                                                                                                                                                                                                                                                        |
| 4    | Modify the /etc/iscsi/iscsid.conf file as described in “Editing the host’s configuration file” in the <i>Setup Guide</i> .                                                                                                                                                                                                                                                                                                                                                        |
| 5    | Ensure that the multipathd service is started before the iscsid service.<br><br>This would require modifying the /etc/init.d/multipathd script and changing the start priority from SXX to S06. This assumes that the default iSCSI installation sets iscsid start priority at S07. If not, you need to ensure that multipathd's start priority is less than iscsid's start priority. The line should be specified as shown in the following example:<br><br># chkconfig: - 06 87 |
| 6    | Run <b>chkconfig multipathd off</b> followed by <b>chkconfig multipathd on</b> .                                                                                                                                                                                                                                                                                                                                                                                                  |
| 7    | Edit the /etc/fstab file and add _netdev mount option for root device. The line should be specified as shown in the following example:<br><br>/dev/mapper/mpath0p1 / ext3 _netdev,defaults 1 1                                                                                                                                                                                                                                                                                    |
| 8    | Reboot the system.                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

**iSCSI node names**

The iSCSI Linux initiator does not support node names with more than 255 characters.

If you use `iscsi nodename` to change the default node name of the storage system, ensure that the new node name does not exceed 255 characters.

## Known problems and limitations using RHEL 4 series

---

### GFS growth is not possible on mounted file system on an iSCSI LUN

**Problem:** GFS provides the online grow feature, where the GFS file system can be extended without unmounting the file system by using the **gfs\_grow** command. This feature cannot be used on a file system created over an iSCSI and multipath device on a NetApp LUN.

---

#### Note

When GFS is used on the LVM volume, it can be extended without unmounting the file system by adding new LUNs to the volume group, and expanding the logical volume on which GFS is created.

---

Though storage systems provide dynamic resizing of the LUN, the iSCSI layer in Linux is not capable of detecting the LUN resize. The new size of the LUN gets reflected on Linux only after the iSCSI service is restarted on Linux.

**Workaround:** Complete the following steps on all the nodes.

| Step | Action                                                       |
|------|--------------------------------------------------------------|
| 1    | Unmount the file system.                                     |
| 2    | Resize the LUN on the storage system.                        |
| 3    | Run the following command:<br><b>multipath -F</b>            |
| 4    | Run the following command:<br><b>service multipath stop</b>  |
| 5    | Run the following command:<br><b>service iscsi restart</b>   |
| 6    | Run the following command:<br><b>service multipath start</b> |
| 7    | Mount the file system.                                       |

| Step | Action                                                                                                   |
|------|----------------------------------------------------------------------------------------------------------|
| 8    | Run the following command on one of the nodes:<br><b>gfs_grow</b>                                        |
| 9    | Run the following command on all nodes to verify the new size:<br><b>running gfs_grow -Tv mountpoint</b> |

For more information, see bug 215923 on Bugs Online at:

<http://now.netapp.com/NOW/cgi-bin/bol>

### Service clvmd restart fails with connect fail on local socket

**Problem:** GFS and the underlying cluster LVM layer are used for providing cluster volume manager. In rare cases, the volume group does not come up on boot. This might happen if iSCSI LUNs are not available before the LVM start-up script comes up.

In such cases, you might have to run `vgscan` after the boot process to read the volume groups from the iSCSI LUNs. This might also need a `clvmd` restart. If you run **service clvmd restart** to restart `clvmd`, in GFS 4.0, it fails. This failure is due to a timing issue in the `clvmd` service start script.

**Workaround:** To avoid the `clvmd` restart failure is to stop and start the `clvmd` by completing the following steps.

| Step | Action                                                                                |
|------|---------------------------------------------------------------------------------------|
| 1    | Run the following command to stop <code>clvmd</code> :<br><b>service clvmd stop</b>   |
| 2    | Run the following command to start <code>clvmd</code> :<br><b>service clvmd start</b> |

## **OCFS2 self-fencing during cluster failover on RHEL 4 Update 3**

**Problem:** During a clustered failover or network disruption, nodes fence themselves by panicking.

Oracle Cluster File System (OCFS2) has a fencing mechanism that relies on each node being able to write a heartbeat value to a specific area on each OCFS2 volume. Every two seconds, each node writes a new heartbeat value, and reads the value of the other node. If a node is unable to write or read the heartbeat value, it retries the operation a specific number of times, and then self-fences by panicking the Linux kernel as it fails to obtain quorum.

During a clustered failover or network disruption, all paths to the storage system can become unavailable, and none of the nodes in the OCFS2 cluster is able to write or read the heartbeat values. As a result, the nodes might fence themselves by panicking.

**Workaround:** The workaround for this problem is to set the `O2CB_HEARTBEAT_THRESHOLD` value in `/etc/sysconfig/o2cb` to a value large enough to ensure that the nodes do not self-fence.

With a dedicated TCP/IP network for the iSCSI link between Linux hosts and storage systems, a reasonable value usually is 61. This corresponds to a total of 120 seconds before a node gives up, and self-fences.

Depending on the load in the iSCSI network used between the Linux host and the storage system, the value should be raised. A value in the range of 61 to 201 would be ideal for a configuration with OCFS2 on mutipath and iSCSI devices.

For more information, see:

[http://oss.oracle.com/projects/ocfs2/dist/documentation/ocfs2\\_faq.html#HEARTBEAT](http://oss.oracle.com/projects/ocfs2/dist/documentation/ocfs2_faq.html#HEARTBEAT)

## System stops responding when shutting down or rebooting

**Problem:** The Linux system stops responding when shutting down or rebooting. The `/etc/init.d/iscsi stop` command might also report errors. This error occurs only in systems running dm-multipath with iSCSI.

**Workaround:** This error is a known problem with the shutdown sequence in RHEL 4 Update 3. If the system is currently not responding, disconnect power to force the shutdown to complete.

To prevent shutdown problems, complete the following steps before shutting down or rebooting.

| Step | Action                                                                                               |
|------|------------------------------------------------------------------------------------------------------|
| 1.   | Unmount all file systems on iSCSI LUNs.                                                              |
| 2.   | Flush the multipath maps by using the following command:<br><b><code>multipath -F</code></b>         |
| 3.   | Stop the iSCSI service by using the following command:<br><b><code>/etc/init.d/iscsi stop</code></b> |
| 4.   | Enter the <b><code>shutdown</code></b> or <b><code>reboot</code></b> command.                        |

For more information, see Red Hat Bugzilla 177355 at:

<https://bugzilla.redhat.com/bugzilla/>

See bug number 183045 on Bugs Online at:

<http://now.netapp.com/NOW/cgi-bin/bol>

## Multiple devices shown for each LUN

**Problem:** The `iscsi-ls -l` and `sanlun lun show` commands display multiple devices for each LUN that is mapped to the Linux host. This problem occurs only with RHEL 4 Update 2 and RHEL 4 Update 3 without dm-multipath.

**Workaround:** The iSCSI initiator creates a separate iSCSI session to each iSCSI portal discovered on the storage system. If you are running RHEL 4 Update 3, you can use the dm-multipath support to manage multiple paths.

Otherwise, configure the storage system or network routing so that only one portal (interface) is visible to the Linux host. See the Host Utilities *Setup Guide* for more information.

## Always map first LUN as LUN 0

**Problem:** If the first LUN mapped to a host has an ID other than 0, the `iscsi-ls -l` command shows incorrect information. Adding LUN 0 later does not solve the problem.

**Workaround:** Always mapping the first LUN for a host to LUN 0 is the default behavior of Data ONTAP when you do not explicitly specify the LUN identifier.

For more information, see NetApp bug 164689 at:

<http://now.netapp.com/NOW/cgi-bin/bol?Type=Detail&Display=164689>.

## Using partitions on Linux with Data ONTAP LUNs might require alignment for best performance

**Problem:** Using default partition settings on LUNs might result in performance problems. This problem can occur if the `sectors` value for a partition is not a multiple of 8.

**Workaround:** Manually enter the disk geometry parameters when creating partitions with the `fdisk` command. For specific steps, see bug number 156121 on Bugs Online at:

<http://now.netapp.com/NOW/cgi-bin/bol>

## iSCSI node names

The iSCSI Linux initiator does not support node names with more than 255 characters.

If you use `iscsi nodename` to change the default node name of the storage system, ensure the new node name does not exceed 255 characters.

## LUNs not recognized by dm-multipath

**Problem:** LUNs are mapped to the Linux host and are correctly reported by the `iscsi-ls -l` and `sanlun lun show all` commands. However, the LUNs are not reported by the `multipath -ll` command. Multipath devices are not created for the LUNs.

**Workaround:** Ensure that the LUNs are not excluded (blacklisted) from the dm-multipath configuration. Enter the following command on the Linux host to display the devices that are blacklisted:

```
multipath -v -d | grep blacklist
```

If any iSCSI devices are blacklisted, check the `devnode_blacklist` section of the `/etc/multipath.conf` file. Pay special attention to any statements added to blacklist local SCSI disks.

---

**Note**

---

devnode "sd[a-b]" excludes not only /dev/sda and /dev/sdb, but also /dev/sdaa, /dev/sdba, and so on. Ensure that you use devnode "sd[a-b]\$" instead.

---

If the iSCSI devices are not blacklisted, but are still not recognized by the multipath command, regenerate the multipath maps by entering the following command:

**multipath -v3**

For more information, see bug number 201302 on Bugs Online at:

<http://now.netapp.com/NOW/cgi-bin/bol>

## **Storage system console displays continuous ISID Rule messages**

**Problem:** The storage system console displays continuous ISID Rule messages, similar to the following:

```
Wed Jun 11 13:42:47 GMT [ss1: iscsi.notice:notice]: ISCSI: iswta,
ISID Rule: new connection from same initiator, shutting down old
session 0
```

```
Wed Jun 11 13:42:48 GMT [ss1: iscsi.notice:notice]: ISCSI: iswta,
ISID Rule: new connection from same initiator, shutting down old
session 1
```

```
Wed Jun 11 13:42:48 GMT [ss1: iscsi.notice:notice]: ISCSI: iswta,
New session from initiator ign.1987-05.com.cisco:rachman at IP addr
10.10.10.10
```

```
Wed Jun 11 13:42:50 GMT [ss1: iscsi.notice:notice]: ISCSI: iswta,
ISID Rule: new connection from same initiator, shutting down old
session 0
```

```
Wed Jun 11 13:42:50 GMT [ss1: iscsi.notice:notice]: ISCSI: iswta,
New session from initiator ign.1987-05.com.cisco:rachman at IP addr
10.10.10.10
```

```
Wed Jun 11 13:42:51 GMT [ss1: iscsi.notice:notice]: ISCSI: iswta,
ISID Rule: new connection from same initiator, shutting down old
```

**Workaround:** Add the following line to the /etc/iscsi.conf file. Ensure that this line is at the beginning of the file, before any "DiscoveryAddress" entries.

```
Continuous=no
```

Then restart the iSCSI service by entering the following command:

**/etc/init.d/iscsi restart**

## Known problems and limitations using RHEL 3 series

---

### Using partitions on Linux with Data ONTAP LUNs might require alignment for best performance

**Problem:** Using default partition settings on LUNs might result in performance problems. This problem can occur if the `sectors` value for a partition is not a multiple of 8.

**Workaround:** Manually enter the disk geometry parameters when creating partitions with the **fdisk** command. For specific steps, see bug number 156121 on Bugs Online at:

<http://now.netapp.com/NOW/cgi-bin/bol>

### Host system performance problems under specific configurations

NetApp testing has found problems with the Linux iSCSI software initiator running under Red Hat Advanced Server (AS) 3.0. The virtual memory (vm) subsystem of the Red Hat AS 3.0 Linux kernels does not interact well with the iSCSI software initiator under certain circumstances. NetApp testing has shown that certain memory and workload configurations might lead to Linux server pauses, and possibly lock up the Linux server completely.

For more information and workarounds, see bug number 149391 at:

<http://now.netapp.com/NOW/cgi-bin/bol?Type=Detail&Display=149391>

### File systems on iSCSI LUNs do not automount after reboot

**Problem:** Red Hat Linux 3 update 4 does not have a standard mechanism for automounting file systems based on iSCSI LUNs after a reboot.

iSCSI LUNs depend on networking services that are typically not available when the file system mounting takes place. Also, the file systems must be mounted by a persistent mechanism; iSCSI devices are assigned device names (`/dev/sda`, `/dev/sdb`...) based on when they are discovered. These device names cannot be guaranteed to map to the same LUN each time.

**Workaround:** You can work around this limitation by manually issuing the **mount -a** command after networking services become available. You could also create a script that waits for networking services to be available before issuing the **mount -a** command.

For more information, see Red Hat Bugzilla 143151 at:

[https://bugzilla.redhat.com/bugzilla/show\\_bug.cgi?id=143151](https://bugzilla.redhat.com/bugzilla/show_bug.cgi?id=143151)

## Unmount iSCSI LUNs before Linux shutdown

**Problem:** File systems on iSCSI LUNs are not automatically unmounted when the Linux system is shut down.

The `iscsi-umountall` script used in earlier versions of the iSCSI initiator is not available with Red Hat 3.0 Update 4. There is no default mechanism for automatically unmounting file systems on iSCSI LUNs when the Linux system is shut down. Shutting down Linux with the file systems still mounted can cause a panic of the Linux host.

**Workaround:** Issue the `umount` command for each file system on an iSCSI LUN before shutting down the Linux host.

For more information, see Red Hat Bugzilla 144781 at:

[https://bugzilla.redhat.com/bugzilla/show\\_bug.cgi?id=144781](https://bugzilla.redhat.com/bugzilla/show_bug.cgi?id=144781)

## Error message: iSCSI: ioctl size 2596 incorrect, expecting 2316

**Problem:** The error message `iSCSI: ioctl size 2596 incorrect, expecting 2316` is displayed on the Linux host.

**Workaround:** Remove the SourceForge iSCSI initiator that was installed for an earlier release of Red Hat Linux. The iSCSI initiator is included in the Red Hat Linux version 3 update 4 and version 4 update 2 kernel and distribution. Do not download an iSCSI initiator from SourceForge.

See the Host Utilities *Setup Guide* for more information.

## sanlun does not display all iSCSI LUNs

**Problem:** The **sanlun** command does not display all the iSCSI LUNs mapped to the initiator.

There are two possible explanations for **sanlun** not displaying all the iSCSI LUNs:

- ◆ Number of SCSI devices

The **sanlun** command checks the number of `/dev/sg*` entries on the host. The default number of `/dev/sg*` entries is 32. If you have more than 32 SCSI devices (for example, two physical SCSI hard drives, one SCSI CD-ROM, and 30 LUNs), **sanlun** does not list all the LUNs.

- ◆ Storage system is unreachable for a long time

If the time during which the storage system is unreachable exceeds the timeouts in the `/etc/iscsi.conf` file, the **sanlun lun show all** command might not provide a complete list of mapped LUNs the first time you run the command.

**Workaround:** The following table lists the workaround for this problem.

| If...                              | Then...                                                                                                                                                                                                                              |
|------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| You have more than 32 SCSI devices | Create more <code>/dev/sg*</code> entries by using the Linux <code>mknod</code> command. For details about the <b>mknod</b> command, see the Linux man page.<br><br><b>Example:</b><br><pre>mknod --mode=660 /dev/sg32 c 21 32</pre> |
| The storage system was unreachable | <ol style="list-style-type: none"><li>1. Run the <b>sanlun lun show all</b> command a second time.</li><li>2. Ensure that the LUNs listed match the devices listed in the output of the <b>iscsi-ls -l</b> command.</li></ol>        |

### **iscsi-ls does not display all iSCSI LUNs**

**Problem:** The **iscsi-ls -l** command lists detailed information about iSCSI LUNs that are mapped to `/dev/sda` through `/dev/sdz`. LUNs mapped to device names starting with `/dev/sdaa` are listed, but most of the information (including the device name) is missing.

**iscsi-ls** not displaying all iSCSI LUNs is a known bug in the **iscsi-ls** command code.

**Workaround:** Use the **sanlun** command to view information about LUNs mapped to iSCSI devices. For example: **sanlun lun show all**.

For more information, see Red Hat Bugzilla 145141 at:

[https://bugzilla.redhat.com/bugzilla/show\\_bug.cgi?id=145141](https://bugzilla.redhat.com/bugzilla/show_bug.cgi?id=145141).

### **Version numbers of iSCSI kernel module and tools do not match**

**Problem:** The iSCSI initiator kernel module displays as 3.6.1, while the version number of the iSCSI daemon and userspace tools displays as 3.6.2.

The daemon and userspace tools were changed late in the release cycle and the kernel module version was not updated to match. The kernel module version 3.6.1 is identical to kernel module 3.6.2.

Version numbers of iSCSi kernel module and tools not matching is not a problem. Therefore, no action is needed.

**File system errors  
and system  
shutdown stops  
responding**

**Problem:** If the initiator is accessing LUNs on a storage system that is unreachable (for example, during a cluster takeover or giveback), file system errors or system shutdown interruption might occur.

File system errors and system shutdown interruption depend on disk command timeout and connection failure timeout settings in the `/etc/iscsi.conf` file.

By default, the disk command and connection failure timeouts are not listed in the `/etc/iscsi.conf` file. As a result, these timeout values are infinite and the host infinitely retries these I/O commands. The infinite timeouts prevent file system errors, but stopping the iSCSI service or a shutdown hangs if the storage system becomes unreachable.

If you did not create Linux file systems on your LUNs, and are using raw devices instead, you can set the timeouts to a value of 180 seconds to prevent system shutdown interruption. However, any application using these LUNs might receive SCSI I/O errors and file system errors if a storage system is unreachable for a long time.

To set the I/O command timeout values, complete the following steps.

**Caution** \_\_\_\_\_  
Do not change the timeout values in the following section if you are using any Linux file systems on your storage system LUNs. This includes ext2 and ext3 file systems.

| Step | Action                                                                                                                                                                                                                                           |
|------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | Ensure that you are logged in to the host with root privileges.                                                                                                                                                                                  |
| 2    | With a text editor, open the host's <code>/etc/iscsi.conf</code> file.                                                                                                                                                                           |
| 3    | At the beginning the file, add the following lines.<br><br>Ensure that you add these lines before the storage system's <code>DiscoveryAddress</code> settings.<br><br><code>MaxDiskCommandTimeout=180</code><br><code>ConnFailTimeout=180</code> |

**iSCSI node names**

The iSCSI Linux initiator does not support node names with more than 255 characters.

If you use `iscsi nodename` to change the default node name of the storage system, ensure that the new node name does not exceed 255 characters.

### **Storage system console displays continuous ISID Rule messages**

**Problem:** The storage system console displays continuous ISID Rule messages, similar to the following:

```
Wed Jun 11 13:42:47 GMT [ss1: iscsi.notice:notice]: ISCSI: iswta,
ISID Rule: new connection from same initiator, shutting down old
session 0
```

```
Wed Jun 11 13:42:48 GMT [ss1: iscsi.notice:notice]: ISCSI: iswta,
ISID Rule: new connection from same initiator, shutting down old
session 1
```

```
Wed Jun 11 13:42:48 GMT [ss1: iscsi.notice:notice]: ISCSI: iswta,
New session from initiator ign.1987-05.com.cisco:rachman at IP addr
10.10.10.10
```

```
Wed Jun 11 13:42:50 GMT [ss1: iscsi.notice:notice]: ISCSI: iswta,
ISID Rule: new connection from same initiator, shutting down old
session 0
```

```
Wed Jun 11 13:42:50 GMT [ss1: iscsi.notice:notice]: ISCSI: iswta,
New session from initiator ign.1987-05.com.cisco:rachman at IP addr
10.10.10.10
```

```
Wed Jun 11 13:42:51 GMT [ss1: iscsi.notice:notice]: ISCSI: iswta,
ISID Rule: new connection from same initiator, shutting down old
```

**Workaround:** Add the following line to the `/etc/iscsi.conf` file. Ensure that this line is at the beginning of the file, before any “DiscoveryAddress” entries.

```
Continuous=no
```

Then restart the iSCSI service by entering the following command:

```
/etc/init.d/iscsi restart
```

# Upgrading the operating system from the previous version

---

## Overview of upgrade procedures

If you are running a version of the operating system, and are upgrading to a later version, unmount the iSCSI LUNs and stop the iSCSI service before you upgrade.

For general information about SUSE Linux Enterprise Server 10 SP2, see

<http://www.novell.com/linux/sp2highlights.html>.

For general information about SUSE Linux Enterprise Server 10 SP1, see

<http://www.novell.com/linux/sp1highlights.html>.

For general information about SUSE Linux Enterprise Server 9 SP2, see

<http://www.novell.com/products/server9/sp2.html>.

For general information about SUSE Linux Enterprise Server 9 SP3, see

<http://www.novell.com/products/server9/sp3.html>.

## Before upgrading to SUSE Linux Enterprise Server 9 SP2 and SP3

Before upgrading the operating system to SUSE Linux Enterprise Server 9 SP2 and SP3, complete the following steps.

| Step | Action                                                                                                                                           |
|------|--------------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | Unmount all existing iSCSI LUNs by running the following command on the SUSE Linux Enterprise Server host:<br><code>/sbin/iscsi-umountall</code> |
| 2    | Stop the iSCSI service on the SUSE Linux Enterprise Server host:<br><code>/etc/init.d/iscsi stop</code>                                          |
| 3    | Continue with the upgrade.                                                                                                                       |

# Known problems and limitations on SUSE Linux Enterprise Server 10 SP1 and SP2

---

## OCFS2 self-fencing during cluster failover

**Problem:** During a clustered failover or network disruption, nodes fence themselves by panicking.

Oracle Cluster File System (OCFS2) has a fencing mechanism that relies on each node being able to write a heartbeat value to a specific area on each OCFS2 volume. Every two seconds, each node writes a new heartbeat value, and reads the value of the other node. If a node is unable to write or read the heartbeat value, it retries the operation a specific number of times, and then self-fences by panicking the Linux kernel as it fails to obtain quorum.

During a clustered failover or network disruption, all paths to the storage system can become unavailable, and none of the nodes in the OCFS2 cluster is able to write or read the heartbeat values. As a result, the nodes might fence by panicking.

**Workaround:** The workaround for this problem is to set the `O2CB_HEARTBEAT_THRESHOLD` value in `/etc/sysconfig/o2cb` to a value large enough to ensure that the nodes do not self-fence.

With a dedicated TCP/IP network for the iSCSI link between Linux hosts and the storage system, a reasonable value usually is 181. This value corresponds to a total of 360 seconds before a node self-fences.

Depending on the load in the iSCSI network, the value should be raised. It is best to use a value in the range of 61 to 201, which would be ideal for a configuration with OCFS2 on multipath, and iSCSI devices.

For more information, see the following:

[http://oss.oracle.com/projects/ocfs2/dist/documentation/ocfs2\\_faq.html#HEARTBEAT](http://oss.oracle.com/projects/ocfs2/dist/documentation/ocfs2_faq.html#HEARTBEAT)

---

### Note

OCFS2 is not supported in SUSE Linux Enterprise Server 10 SP2.

---

## LUN status undefined

**Problem:** `multipath -l` displays the LUN status as `undef`.

**Workaround:** Use `multipath -ll` to get the correct information.

## LUNs not recognized by dm-multipath

**Problem:** LUNs are mapped to the Linux host and are correctly reported by the `iscsi-ls -l` and `sanlun lun show all` commands. However, the LUNs are not reported by the `multipath -ll` command. Multipath devices are not created for the LUNs.

**Workaround:** Ensure that the LUNs are not excluded (blacklisted) from the dm-multipath configuration. Enter the following command on the Linux host to display the devices that are blacklisted:

```
multipath -v3 -d | grep blacklist
```

If any iSCSI devices are blacklisted, check the `blacklist` section of the `/etc/multipath.conf` file. Ensure that the entries are correctly specified.

If the iSCSI devices are not blacklisted, but are still not recognized by the multipath command, regenerate the multipath maps by entering the following command:

```
multipath -v3
```

## iSCSI node names

The iSCSI Linux initiator does not support node names with more than 255 characters.

## File system errors and system shutdown stops responding

**Problem:** If the initiator is accessing LUNs on a storage system and I/O takes a long time (for example, the target is unreachable, a CFO is in progress, or the target is heavily loaded), file system errors or system shutdown interruption might occur.

## File systems on partitions might have performance problems

**Problem:** The default `fdisk` command parameters used to create a partition on a LUN might result in poor performance for that LUN.

The default `fdisk` parameters can create a misalignment between the LUN block boundaries and the underlying WAFL® boundaries. This misalignment causes the storage system to double its I/O workload for that LUN.

**Workaround:** You can explicitly specify the `fdisk` parameters when creating a partition to use values that guarantee the correct block alignment.

For specific steps, see bug number 156121 on Bugs Online at:

<http://now.netapp.com/NOW/cgi-bin/bol>

**Online resize of file system is not supported**

**Problem:** Currently, a resize of the LUN on the NetApp storage controller is reflected on the host by rescanning the iSCSI sessions. The rescan reflects the new size of the LUN. However, the resize of the file system is not reflected while it is online. See the Novell bugzilla number 386325 for further details.

**Workaround:** To reflect this new size of the LUN on the file system, unmount the file system and use the resize utility to extend or shrink the file system.

# Known problems and limitations on SUSE Linux Enterprise Server 9

---

## OCFS2 self-fencing during cluster failover on SUSE Linux Enterprise Server 9 SP3

**Problem:** During a clustered failover or network disruption, nodes fence themselves by panicking.

Oracle Cluster File System (OCFS2) has a fencing mechanism that relies on each node being able to write a heartbeat value to a specific area on each OCFS2 volume. Every two seconds, each node writes a new heartbeat value, and reads the value of the other node. If a node is unable to write or read the heartbeat value, it retries the operation a specific number of times, and then self-fences by panicking the Linux kernel as it fails to obtain quorum.

During a clustered failover or network disruption, all paths to the storage system can become unavailable, and none of the nodes in the OCFS2 cluster is able to write or read the heartbeat values. As a result, the nodes might fence themselves by panicking.

**Workaround:** The workaround for this problem is to set the `O2CB_HEARTBEAT_THRESHOLD` value in `/etc/sysconfig/o2cb` to a value large enough to ensure that the nodes do not self-fence.

With a dedicated TCP/IP network for the iSCSI link between Linux hosts and storage system, a reasonable value usually is 61. This corresponds to a total of 120 seconds before a node gives up, and self-fences.

Depending on the load in the iSCSI network used between the Linux host and the storage system, the value should be raised. A value in the range of 61 to 201 would be ideal for a configuration with OCFS2 on mutipath, and iSCSI devices.

For more information, see:

[http://oss.oracle.com/projects/ocfs2/dist/documentation/ocfs2\\_faq.html#HEARTBEAT](http://oss.oracle.com/projects/ocfs2/dist/documentation/ocfs2_faq.html#HEARTBEAT)

See also, bug 210340 on Bugs Online at:

<http://now.netapp.com/NOW/cgi-bin/bol>

**Automatically mounting and automatically unmounting the file system on multipath devices do not work on SUSE Linux Enterprise Server 9 SP3**

**Problem:** Automated mounting of multipath-based file systems at boot time using entries in `/etc/fstab.conf` or `/etc/fstab.iscsi` does not work. This is because the placement of the multipath boot and shutdown scripts in the rc script list in SUSE Linux Enterprise Server 9 SP3 is out of order with respect to the mount and unmount scripts for iSCSI and non-iSCSI devices.

The multipathd and iscsid startup scripts are at the same serial number for a run level in `/etc/init.d/rc3.d`. Because alphabetically *i* comes before *m*, iscsid starts before multipathd.

As a result, the multipath boot and shutdown sequence is not synchronized with the iSCSI boot and shutdown sequence. Multipath devices are not available when the mount command is run, and the mount command fails to mount these devices.

Similarly in the shutdown path, by the time the unmount script runs, the multipath utility would have already been shut down.

**Workaround:** Manual mounting and unmounting is required for multipath-based file systems on SUSE Linux Enterprise Server 9 SP3. After the multipath devices are available on booting, the file systems can be mounted.

For unmounting, ensure that any multipath-based file systems are unmounted before the multipath utility is shut down.

To ensure the proper shutdown of multipath-based iSCSI file systems, complete the following steps.

| Step | Action                                                                                                          |
|------|-----------------------------------------------------------------------------------------------------------------|
| 1    | Unmount all file systems on iSCSI LUNs.                                                                         |
| 2    | Flush the multipath maps by entering the following command:<br><b><code>multipath -F</code></b>                 |
| 3    | Stop the multipath daemon by entering the following command:<br><b><code>/etc/init.d/multipathd stop</code></b> |
| 4    | Stop the iSCSI service by entering the following command:<br><b><code>/etc/init.d/iscsi stop</code></b>         |

For more information, see bug 196997 on Bugs Online at:

<http://now.netapp.com/NOW/cgi-bin/bol>.

### LUNs not recognized by dm-multipath

**Problem:** LUNs are mapped to the Linux host and are correctly reported by the `iscsi-ls -l` and `sanlun lun show all` commands. However, the LUNs are not reported by the `multipath -ll` command. Multipath devices are not created for the LUNs.

**Workaround:** Ensure that the LUNs are not excluded (blacklisted) from the dm-multipath configuration. Enter the following command on the Linux host to display the devices that are blacklisted:

```
multipath -v3 -d | grep blacklist
```

If any iSCSI devices are blacklisted, check the `blacklist` section of the `/etc/multipath.conf` file. Ensure that the entries are correctly specified.

If the iSCSI devices are not blacklisted, but are still not recognized by the multipath command, regenerate the multipath maps by entering the following command:

```
multipath -v3
```

### File systems on partitions might have performance problems

**Problem:** The default `fdisk` command parameters used to create a partition on a LUN might result in poor performance for that LUN.

The default `fdisk` parameters can create a misalignment between the LUN block boundaries and the underlying WAFL boundaries. This misalignment causes the storage system to double its I/O workload for that LUN.

**Workaround:** You can explicitly specify the `fdisk` parameters when creating a partition to use values that guarantee the correct block alignment. For specific steps, see bug number 156121 on Bugs Online at:

<http://now.netapp.com/NOW/cgi-bin/bol>

### iscsi-ls command displays error message

**Problem:** For some SUSE Linux versions, the `iscsi-ls -l` command displays the following error message for each LUN:  
Cannot open /etc/sysconfig/devlabel: No such file or directory

This message is displayed because of a bug in the command. The message does not indicate an error and can be safely ignored.

## One or more ext3 file systems not mounted after reboot

**Problem:** After a reboot of the host, one or more ext3 file systems are not mounted, and the following error message is displayed:  
fs type ext3 not supported by kernel

The module that supports ext3 file systems was not loaded before the system tried to mount the file systems.

**Workaround:** Issue the **iscsi-mountall** command to mount the file systems.

## sanlun does not display all iSCSI LUNs

**Problem:** The **sanlun** command does not display all iSCSI LUNs mapped to the initiator.

There are two possible explanations for **sanlun** not displaying all iSCSI LUNs:

- ◆ Number of SCSI devices

The **sanlun** command checks the number of `/dev/sg*` entries on the host. The default number of `/dev/sg*` entries is 32. If you have more than 32 SCSI devices (for example, two physical SCSI hard drives, one SCSI CD-ROM, and 30 LUNs), **sanlun** does not list all the LUNs.

- ◆ Storage system is unreachable for a long time

If the time during which the storage system is unreachable exceeds the timeouts in the `/etc/iscsi.conf` file, the **sanlun lun show all** command might not provide a complete list of mapped LUNs the first time you run the command.

**Workaround:** The following table lists the workaround for this problem.

| If...                              | Then...                                                                                                                                                                                                                                               |
|------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| You have more than 32 SCSI devices | <p>Create more <code>/dev/sg*</code> entries by using the Linux <code>mknod</code> command.</p> <p>For details about the <code>mknod</code> command, see the Linux man page.</p> <p><b>Example:</b></p> <pre>mknod --mode=660 /dev/sg32 c 21 32</pre> |

| If...                              | Then...                                                                                                                                                                                                                                                      |
|------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| The storage system was unreachable | <ol style="list-style-type: none"> <li>1. Run the <b>sanlun lun show all</b> command a second time.</li> <li>2. Ensure that the LUNs listed match the devices listed in the /dev/iscsi directory or in the output of the <b>iscsi-ls</b> command.</li> </ol> |

## File system errors and system shutdown stops responding

**Problem:** If the initiator is accessing LUNs on a storage system and I/O takes a long time (for example, the target is unreachable, a CFO is in progress, or the target is heavily loaded), file system errors or system shutdown interruption might occur.

File system errors and system shutdown interruption depend on disk command timeout and connection failure timeout settings in the /etc/iscsi.conf file.

By default, the disk command and connection failure timeouts are not listed in the /etc/iscsi.conf file. As a result, these timeout values are infinite and the host infinitely retries these I/O commands. The infinite timeouts prevent file system errors, but stopping the iSCSI service hangs or a shutdown is interrupted if storage system I/O takes a long time.

If you did not create Linux file systems on your LUNs, and are using raw devices instead, you can set the timeouts to a value of 180 seconds to prevent system shutdown interruption. However, any application using these LUNs might receive SCSI I/O errors and file system errors if a storage system is unreachable for a long time.

To set the I/O command timeout values, complete the following steps:

### Caution

Do not change the timeout values in the following section if you are using any Linux file systems on your storage system LUNs. This includes Reiser, ext2, and ext3 file systems.

| Step | Action                                                          |
|------|-----------------------------------------------------------------|
| 1    | Ensure that you are logged in to the host with root privileges. |
| 2    | With a text editor, open the host's /etc/iscsi.conf file.       |

| Step | Action                                                                                                                                                                                                     |
|------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 3    | <p>At the beginning the file, add the following lines. Ensure that you add these lines before the storage system's DiscoveryAddress settings:</p> <pre>MaxDiskCommandTimeout=180 ConnFailTimeout=180</pre> |

## iSCSI node names

The iSCSI Linux initiator does not support node names with more than 255 characters.

If you use `iscsi nodename` to change the default node name of the storage system, ensure that the new node name does not exceed 255 characters.

## Information about additional problems

---

Occasionally, additional problems are reported with Host Utilities, or known bugs are fixed. The problems are listed in an online system called Bugs Online.

For a list of known and fixed bugs, access Bugs Online by completing the following steps.

| Step | Action                                                                                                            |
|------|-------------------------------------------------------------------------------------------------------------------|
| 1    | Log in to <a href="http://now.netapp.com">http://now.netapp.com</a> .                                             |
| 2    | Click Service and Support.                                                                                        |
| 3    | Click Bugs Online and Release Tools.                                                                              |
| 4    | In the Bug Search Tools section, select iSCSI-Linux from the Bug Type(s) menu.                                    |
| 5    | Click Go!.<br><br><b>Result:</b> Bugs Online - Advanced Search Results lists all bugs for Linux iSCSI initiators. |

## Where to find additional information

### Where to go for more information

The following table describes documents with the most current information about host and storage system requirements and additional documentation.

| If you want...                                                                        | Go to...                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| The most current system requirements for the Linux host                               | <p>The NetApp Interoperability Matrix and NetApp SAN Support Matrix at the following sites:</p> <ul style="list-style-type: none"><li>◆ <a href="http://now.netapp.com/NOW/products/interoperability">http://now.netapp.com/NOW/products/interoperability</a></li><li>◆ <a href="http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/index.shtml">http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/index.shtml</a>.</li></ul>                                                                                                                                                                    |
| The supported storage system models for Data ONTAP licensed with iSCSI                | <ul style="list-style-type: none"><li>◆ The NetApp Interoperability Matrix and NetApp SAN Support Matrix at the following sites:<ul style="list-style-type: none"><li>❖ <a href="http://now.netapp.com/NOW/products/interoperability">http://now.netapp.com/NOW/products/interoperability</a></li><li>❖ <a href="http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/index.shtml">http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/index.shtml</a>.</li></ul></li><li>◆ <i>System Configuration Guide</i> for Data ONTAP at <a href="http://now.netapp.com">http://now.netapp.com</a>.</li></ul> |
| Information about how to configure and manage the iSCSI service on the storage system | <p>The following documents for your Data ONTAP release:</p> <ul style="list-style-type: none"><li>◆ <i>Data ONTAP Block Access Management Guide for iSCSI</i></li><li>◆ <i>Data ONTAP Release Notes</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                       |
| Best practices or configuration issues                                                | <p>The NetApp Knowledgebase at:</p> <ul style="list-style-type: none"><li>◆ <a href="https://now.netapp.com/eservice/kbbrowse">https://now.netapp.com/eservice/kbbrowse</a></li><li>◆ <a href="http://now.netapp.com/Knowledgebase/solutionarea.asp?id=kb8190">http://now.netapp.com/Knowledgebase/solutionarea.asp?id=kb8190</a></li></ul>                                                                                                                                                                                                                                                                             |
| Information about dm-multipath                                                        | <p>The Multipath example files included with the device-mapper-multipath RPM are in /usr/share/doc/device-mapper-multipath-&lt;version&gt;.</p> <p>&lt;version&gt; is the latest version number shipped with the RHEL release.</p>                                                                                                                                                                                                                                                                                                                                                                                      |
| Information about Red Hat Cluster and GFS                                             | <a href="http://www.redhat.com/docs/manuals/csgfs/">http://www.redhat.com/docs/manuals/csgfs/</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

| If you want...                                                                           | Go to...                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Information about OCFS2                                                                  | <ul style="list-style-type: none"> <li>◆ Oracle Cluster File System (OCFS2) <i>User's Guide</i>:<br/> <a href="http://oss.oracle.com/projects/ocfs2/dist/documentation/ocfs2_users_guide.pdf">http://oss.oracle.com/projects/ocfs2/dist/documentation/ocfs2_users_guide.pdf</a></li> <li>◆ OCFS2 - Frequently Asked Questions:<br/> <a href="http://oss.oracle.com/projects/ocfs2/dist/documentation/ocfs2_faq.html">http://oss.oracle.com/projects/ocfs2/dist/documentation/ocfs2_faq.html</a> (The FAQs on OCFS2)</li> </ul> |
| Information about Oracle RAC                                                             | <a href="http://www.oracle.com/technology/pub/articles/smiley_rac10g_install.html">http://www.oracle.com/technology/pub/articles/smiley_rac10g_install.html</a>                                                                                                                                                                                                                                                                                                                                                                |
| NetApp Knowledge Base for Oracle RAC                                                     | <p>Knowlegde base for Oracle RAC at the following sites:</p> <ul style="list-style-type: none"> <li>◆ <a href="http://www.netapp.com/library/tr/3339.pdf">http://www.netapp.com/library/tr/3339.pdf</a></li> <li>◆ <a href="http://www.netapp.com/library/tr/3423.pdf">http://www.netapp.com/library/tr/3423.pdf</a></li> </ul>                                                                                                                                                                                                |
| Information about configuring multipathing on SUSE Linux Enterprise Server               | <p>The Novell knowledgebase article <i>How to setup / use multipathing on SLES</i> at:<br/> <a href="http://www.novell.com/support">http://www.novell.com/support</a></p>                                                                                                                                                                                                                                                                                                                                                      |
| Information about configuring the iSCSI initiator on SUSE Linux Enterprise Server 10 SP1 | <p>SUSE Linux Enterprise Server 10 SP1 <i>Installation and Administration Guide</i> at:<br/> <a href="http://www.novell.com/documentation/sles10/">http://www.novell.com/documentation/sles10/</a></p>                                                                                                                                                                                                                                                                                                                         |